This article has been accepted for publication in IEEE Transactions on Evolutionary Computation. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TEVC.2026.3749205

Multiobjective Stackelberg Security Games: Problem Formulation, Benchmarks and Evolutionary Algorithms

Adam Żychowski, Abhishek Gupta, *Senior Member, IEEE*, Yew-Soon Ong, *Fellow, IEEE*, and Jacek Mańdziuk, *Senior Member, IEEE*

***Abstract*—This paper introduces a comprehensive framework for Multiobjective Stackelberg Security Games (MO-SSGs), extending traditional single-objective formulations to settings where both the defender (Leader) and attacker (Follower) optimize vector-valued payoffs. We define equilibrium concepts for these games and establish their key properties. A hallmark of the MO-SSG formulation is the explicit use of mixed strategies by the Leader—defined as a probability distribution over pure strategies—to confuse the Follower. Another crucial feature is that a single Leader strategy can yield a *set* of vector-valued payoffs due to the existence of multiple optimal Follower responses. To address this strategy-to-payoff set mapping, we introduce the concept of worst-case dominance, a pessimistic approach for evaluating strategies against adversarial Followers. This formulation precludes direct application of standard Pareto-based Multiobjective Evolutionary Algorithms (MOEAs). Consequently, we propose two specialized population-based search methods: the Worst-Case Dominance Evolutionary Algorithm (WCD-EA) and the Worst-Case Dominance Coevolutionary Algorithm (WCD-CoEA). These algorithms utilize novel evolutionary selection and variation operators to facilitate convergence toward worst-case nondominated strategy sets while preserving diversity in the Leader's multidimensional objective space. We validate our approach using a new benchmark suite of instructional and practically-inspired MO-SSG instances. Experimental results demonstrate the computational feasibility of the framework and characterize the trade-offs between solution quality and efficiency. Ultimately, this work demonstrates that the unique landscape of MO-SSGs calls for the design of new evolutionary methods and operators, offering a foundation for future research into multiobjective sequential defender–attacker models.**

***Index Terms*—Multiobjective Games, Evolutionary Algorithms, Game Theory, Security Management**

Abhishek Gupta was supported by the Ramanujan Fellowship from the Anusandhan National Research Foundation (ANRF), Government of India under Grant RJF/2022/000115. Yew-Soon Ong was partly supported by the MTI under its AI Centre of Excellence for Manufacturing (AIMfg) (Award W25MCMF014), the A*STAR-NTU Research Joint Lab for Smart and Sustainable Advanced Manufacturing, and the College of Computing and Data Science, Nanyang Technological University. Jacek Mańdziuk was partially supported by the National Science Centre, Poland, grant number 2023/49/B/ST6/01404.

Adam Żychowski and Abhishek Gupta contributed equally to this work. Corresponding authors: Adam Żychowski and Abhishek Gupta (e-mail: adam.zychowski@pw.edu.pl; abhishekgupta@iitgoa.ac.in).

A. Żychowski is with the Faculty of Mathematics and Information Science, Warsaw University of Technology, Poland.

A. Gupta is with the School of Mechanical Sciences, Indian Institute of Technology Goa, India.

Y. S. Ong is with the College of Computing and Data Science, Nanyang Technological University, Singapore, and the Agency for Science, Technology and Research, Singapore.

J. Mańdziuk is with the Faculty of Mathematics and Information Science, Warsaw University of Technology, Poland, and the Faculty of Computer Science, AGH University of Krakow, Poland.

I. Introduction

The increasing complexity of modern security challenges, from cyber warfare [1] to wildlife conservation [2], [3], has exposed critical limitations in traditional game-theoretic models that reduce strategic decision-making to single-dimensional utility optimization. Traditional Stackelberg Security Games (SSGs) have been successfully applied in high-impact domains [4], [5] such as airport security [6], [7], ferry protection [8], [9], or poaching prevention [2], [10]. However, these models typically reduce strategic reasoning to a single-objective optimization problem, thereby oversimplifying the rich set of trade-offs faced by real-world decision makers.

In practice, defenders often balance *multiple, conflicting objectives*: e.g., minimizing operational cost while ensuring sufficient coverage, deterring adversaries while preserving public trust, or allocating limited resources while maintaining system resilience. Similarly, attackers may simultaneously seek to maximize impact while minimizing the risk of exposure. Classical single-objective SSGs do not capture the trade-offs among such competing priorities, leading to strategies that may be optimal for one objective yet highly suboptimal for others.

An example of the practical necessity for multiobjective security games emerges from wildlife conservation. Anti-poaching efforts naturally involve multiple competing goals for both defenders and attackers. Park rangers must balance several priorities at once: maximizing wildlife protection, minimizing costs, maintaining good relationships with local communities, and ensuring long-term sustainability. These objectives often conflict with each other—better protection usually means higher costs, while community engagement may require reducing patrol intensity in certain areas. Similarly, poachers face their own competing objectives: maximizing profit while minimizing the risk of being caught, reducing travel costs, and avoiding dangerous confrontations with rangers. The Protection Assistant for Wildlife Security (PAWS) system [11], deployed in Uganda's Queen Elizabeth National Park and other conservation use cases, show how these competing priorities work in practice. A key challenge is that decision-makers often don't know how to weigh these different objectives *a priori* when they need to make strategic decisions. Rangers might be unsure whether saving money is more important than slightly better coverage. Similar to standard multiobjective optimization, this uncertainty about preferences makes traditional approaches that combine objectives into a single score inadequate, since stakeholders first need to understand what trade-offs are possible before deciding which ones they prefer.

Recent work has introduced *multiobjective security games*

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Transactions on Evolutionary Computation. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TEVC.2026.3749205

where the defender's utility is modeled by a vector of objectives [12], [13]. While these approaches provide valuable insights into trade-offs on the defender's side, they retain the simplifying assumption of a single-objective attacker. Consequently, these models fail to capture the essence of adversarial interactions where both sides face complex, heterogeneous and possibly conflicting goals which is a common situation in modern security domains. To the best of our knowledge, no prior work has studied *multiobjective Stackelberg security games* (MO-SSGs) where *both* the defender (Leader) and attacker (Follower) optimize vector-valued payoffs.

This paper addresses this gap by introducing a general theoretical and algorithmic framework for MO-SSGs. We formalize the concept of *rationalizable Leader strategies* for vector-valued security games, enabling principled reasoning about equilibria under postponed preference articulation. A characteristic feature that distinguishes the MO-SSG formulation from previous work on multiobjective games [14], [15] is the explicit use of *mixed strategies*—defined as a probability distribution over pure strategies—to confuse the opponent. Moreover, unlike conventional multiobjective optimization where a candidate solution maps to a single vector-valued payoff, a strategy in an MO-SSG can result in a *set* of vector-valued payoffs, contingent upon the Follower's optimal responses. To ensure robustness against adversarial Followers, we introduce the concept of pessimistic, *worst-case dominance* between candidate strategies. The worst-case dominance structure nullifies direct applicability of existing Pareto dominance-based multiobjective evolutionary algorithms (MOEAs), and the pessimistic formulation precludes the use of EAs crafted for multiobjective bilevel programs that typically assume optimistic Leader-Follower interactions [16], [17]. Consequently, we propose two novel population-based search methods: the Worst-Case Dominance Evolutionary Algorithm (WCD-EA), which performs exhaustive evaluation of Follower responses, and the Worst-Case Dominance Coevolutionary Algorithm (WCD-CoEA), which coevolves Follower strategies. Our approach presents specialized evolutionary selection and variation operators to facilitate convergence towards the worst-case nondominated strategy set while simultaneously promoting diversity in the Leader's multidimensional objective space.

Beyond theoretical and algorithmic novelty, our framework has direct relevance for real-world applications in cyber defense, infrastructure security, environmental protection, and resource allocation under uncertainty. By explicitly accounting for multiple objectives on *both* sides of the security game, MO-SSGs enable robust and interpretable decision-making in complex adversarial environments. The main contributions of this paper can be summarized as follows:

- **Extended theoretical framework.** We introduce the basic definition of Multiobjective Stackelberg Security Games (MO-SSGs), in which both players are characterized by vector-valued payoffs. We describe the resulting equilibrium concepts, including set-based worst-case dominance relations and rationalizable Leader strategies.
- **Benchmark suite for MO-SSGs.** To facilitate systematic evaluation and reproducibility, we design and release a comprehensive benchmark suite covering two distinct security domains. The first introduces an instructional *Tug of War* (ToW) game, a model of physical perimeter defense characterized by continuous spatial strategies. The second extends canonical *FlipIt* games to multiobjective settings, capturing discrete cybersecurity interactions on networks. Together, these instances provide a standardized testbed for future research, encompassing both combinatorial and geometric game structures.
- **Evolutionary algorithms and experimental study.** We develop the WCD-EA and WCD-CoEA tailored to MO-SSGs, enabling efficient exploration of strategy spaces for both Leader and Follower under postponed preference articulation. The worst-case dominance structure necessitates a departure from off-the-shelf solvers: since a single Leader strategy maps to a *set* of payoff vectors, straightforward application of standard MOEAs to this formulation is precluded. Extensive computational experiments on the benchmark problems validate the effectiveness of our newly developed methods, and provide insights into the practical solvability of MO-SSGs.

II. Preliminaries

A. Single-objective Stackelberg Security Games

In the single-objective setting, a Stackelberg Security Game (SSG) generally comprises two players: a Leader (L) and a Follower (F). Let Π_L represent the set of all possible Leader strategies (pure or mixed), Π_F represent the set of all possible Follower strategies, $u_L(\pi_L, \pi_F) \in \mathbb{R}$ denote the Leader's expected payoff, and $u_F(\pi_L, \pi_F) \in \mathbb{R}$ denote the Follower's expected payoff. $\pi_L \in \Pi_L$ and $\pi_F \in \Pi_F$ are specific strategies chosen by the Leader and the Follower, respectively. As discussed further in Section IV-B, SSG formulations typically assume pure Follower strategies. SSGs represent an important class of security games where the Leader commits to a strategy first, and the Follower responds optimally after observing this commitment. In sequential games spanning multiple time steps, both players commit to their respective strategies (defined as complete, pre-planned sequences of actions across all time steps) prior to the actual execution of the game. Consequently, the Follower responds to the Leader's overall strategy commitment as a whole, and neither player dynamically adapts their actions step-by-step based on the opponent's realized moves during execution. The Leader's objective is then to find a strategy that maximizes their expected utility, considering the Follower's best response.

In the single-objective case, the Stackelberg equilibrium is a pair of strategies (π_L^*, π_F^*) that satisfy the condition:

$$\pi_L^* = \arg \max_{\pi_L \in \Pi_L} u_L(\pi_L, \pi_F^*(\pi_L)),$$

where $\pi_F^*(\pi_L)$ is the Follower's best response to π_L, given as:

$$\pi_F^*(\pi_L) = \arg \max_{\pi_F \in \Pi_F} u_F(\pi_L, \pi_F).$$

A key consideration in the SSG formulation is tie-breaking. When multiple Follower strategies yield the same payoff,

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Transactions on Evolutionary Computation. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TEVC.2026.3749205

the *strong* Stackelberg equilibrium assumes that the Follower breaks ties in favor of the Leader [18]. This is justified since the Leader can induce the Follower to choose favorably via an infinitesimal adjustment to their strategy [19].

B. Background Literature

Research on SSGs has produced a wide range of solution methods, modeling extensions, and real-world applications. In the following, we briefly review key contributions in these three directions.

1) Solution Methods: Computational methods for SSGs range from exact mathematical programming to scalable heuristics. Early exact approaches like DOBSS [20] relied on MILP, while subsequent methods improved scalability using sequence-form representations [21] and Correlated Equilibrium concepts (SEFCE) [22]. However, scalability limits in large games motivated heuristic techniques. The O2UCT algorithm [23], [24] adapts Monte Carlo Tree Search [25] for adversarial planning, while evolutionary approaches like EASG [26], [27] and coevolutionary extensions [28] efficiently explore mixed strategy spaces. Recently, explicit promotion of sparsity in these algorithms has produced state-of-the-art results for problems intractable for exact solvers [29], [30].

2) Game Variants: Beyond standard formulations, research has expanded to capture richer behavioral assumptions, such as *signaling games* [10], [31] and *bounded rationality* models (e.g., Quantal Response, Prospect Theory,Anchoring models) that improve robustness against human adversaries [32]–[36]. Domain-specific variants also exist, including *spatio-temporal patrolling games* [37], [38], *green security games* for wildlife protection and conservation [2], [39], and *cybersecurity defense games* [1], [40]. These variants exploit combinatorial structures (such as paths, flows, or networks) to enable scalable solution techniques while reflecting domain constraints.

3) Real-World Applications: SSGs have been successfully deployed in high-impact domains, including ARMOR for airport security [41], PROTECT for port safety [38], and PAWS [11] for combat poaching. These deployments highlight the practical utility and robustness of SSGs against adaptive adversaries and limited resources. The diversity of variants and deployments underscores the maturity and broad applicability of this research area. However, while the single-objective SSG literature is extensive and well-developed, many real-world security scenarios naturally involve multiple, often conflicting objectives. Dedicated research on multiobjective formulations of SSGs remains scarce, which motivates the need for novel approaches that explicitly account for such trade-offs.

III. Formulation and Concepts of MO-SSGs

In this section, we consider SSGs where both the Leader and the Follower have multiple objectives of interest. That is, the payoff functions are vector-valued:

$$\mathbf{u}_L(\pi_L, \pi_F) \in \mathbb{R}^{K_L}, \quad \mathbf{u}_F(\pi_L, \pi_F) \in \mathbb{R}^{K_F},$$

which, considering a mixed strategy by the Leader, represent *expected* payoff vectors calculated with respect to the underlying probability distribution over pure strategies. This setting requires both players to evaluate Pareto-optimal sets of strategies rather than single best responses, significantly increasing complexity. Notice that given a particular strategy (pure or mixed) by the Leader, there could exist a set of responses by the Follower that are all optimal in the multiobjective sense. Hence, a single strategy yields a set of possible payoff vectors, i.e., we get a point-to-set mapping that departs significantly from the point-to-point mappings usually encountered in multiobjective optimization. For example, the objective spaces depicted in Figure 1 (and discussed further in Section III-E) show that a single strategy chosen by the Leader, such as Strategy B, can evoke multiple incomparable Pareto-optimal responses from the Follower, such as responses 2 and 3. This effectively maps the Leader's single decision point to a set of distinct payoff outcomes.

Even in this point-to-set regime, the concept of Pareto dominance provides a basis for developing a principled means of comparing strategies in MO-SSGs. Below, we recall the definition of Pareto dominance, and present the conventional definition of set-based dominance that conforms to it [42].

Definition 1 (Pareto dominance between payoff vectors)**.** *Let $\boldsymbol{f} = [f_1, f_2, \ldots, f_K] \in \mathbb{R}^K$ and $\boldsymbol{h} = [h_1, h_2, \ldots, h_K] \in \mathbb{R}^K$ be two vectors in objective space. Then, $\boldsymbol{f}$ Pareto dominates $\boldsymbol{h}$ in a maximization problem ($\boldsymbol{f} \overset{max}{\succ} \boldsymbol{h}$) if:*

$$f_k \geq h_k \ \forall k \in \{1, \ldots, K\}, \text{ and } \exists k \in \{1, \ldots, K\} \ s.t. \ f_k > h_k.$$

Definition 2 (Dominance between sets of vectors)**.** *Let F and H be sets of vectors in objective space. Set F dominates set H in a maximization problem, denoted $F \overset{max}{\succ} H$, if $\forall \boldsymbol{h} \in H$, $\exists \boldsymbol{f} \in F$ such that $\boldsymbol{f} \overset{max}{\succ} \boldsymbol{h}$.*

These definitions apply for domination relations in minimization ($\overset{\min}{\succ}$) by reversing the inequalities in Definition 1.

A. Optimal Follower Responses in MO-SSGs

For a fixed Leader strategy $\pi_L \in \Pi_L$, the Follower's rational responses are those that are Pareto-optimal given their objectives. Formally, we define this set as:

$$S_F^{\text{rat}}(\pi_L) = \{\pi_F \in \Pi_F : \nexists \pi'_F \in \Pi_F \text{ such that } \mathbf{u}_F(\pi_L, \pi'_F) \overset{\max}{\succ} \mathbf{u}_F(\pi_L, \pi_F)\}.$$

All members of this set are mutually incomparable [43]; consequently, under postponed preference articulation, they are effectively *tied*.

B. Leader's Worst-Case Outcome Set

Without prior information about Follower preferences, the Leader remains uncertain about the strategy the Follower is likely to enact from within the set $S_F^{\text{rat}}(\pi_L)$ of rational Follower responses. The strong Stackelberg equilibrium assumption of breaking ties in favor of the Leader suggests a degree of cooperation and optimism that is deemed unrealistic for adversarial game settings [44]. Although this optimistic condition is widely adopted in the literature due to its mathematical convenience [45], in this paper we take a distinct conservative approach. Specifically, the Leader is

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Transactions on Evolutionary Computation. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TEVC.2026.3749205

assumed to act robustly by evaluating the set of *worst-case* payoffs, considering all Pareto-minimal outcomes among the Follower's response set as:

$$\mathcal{W}_L(\pi_L) = \{\mathbf{u}_L(\pi_L, \pi_F) : \pi_F \in S_F^{\text{rat}}(\pi_L), \nexists \pi'_F \in S_F^{\text{rat}}(\pi_L) \text{ such that } \mathbf{u}_L(\pi_L, \pi_F) \overset{\text{max}}{\succ} \mathbf{u}_L(\pi_L, \pi'_F)\}.$$

$\mathcal{W}_L(\pi_L)$ contains all Leader payoffs that are *minimal* in the Pareto sense among rational Follower responses, representing the Leader's worst-case (i.e., pessimistic) outcomes. Since the binary relation $\overset{\text{max}}{\succ}$ is a partial order that admits incomparable elements, $\mathcal{W}_L(\pi_L)$ is generally not a singleton. Consequently, the Leader's payoffs are structurally point-to-set mappings. This differs from standard set-based multiobjective optimization, where sets are typically collections of independent point-to-point solutions [42]. This structural difference marks a unique feature of the MO-SSG formulation.

C. Worst-Case Dominance Between Leader Strategies

In line with our robust formulation, we define a set-based preference relation to compare two Leader strategies under pessimistic point-to-set evaluations.

Definition 3 (Worst-case dominance relation between sets)**.** *Given two Leader strategies $\pi_L, \pi'_L \in \Pi_L$, the set $\mathcal{W}_L(\pi'_L)$* worst-case dominates *$\mathcal{W}_L(\pi_L)$, i.e., $\mathcal{W}_L(\pi'_L) \overset{\text{w.c.}}{\succ} \mathcal{W}_L(\pi_L)$, in a maximization problem if and only if $\mathcal{W}_L(\pi_L) \overset{\text{min}}{\succ} \mathcal{W}_L(\pi'_L)$.*

This relationship prefers Leader strategies that are robust under rational and adversarial Follower behavior. It is important to note that $\overset{\text{w.c.}}{\succ}$ and conventional set-based dominance (Definition 2) are in general logically independent. Moreover, the logical equivalence $\mathcal{W}_L(\pi'_L) \overset{\text{w.c.}}{\succ} \mathcal{W}_L(\pi_L) \iff \mathcal{W}_L(\pi_L) \overset{\text{min}}{\succ} \mathcal{W}_L(\pi'_L)$ guarantees that, for every objective, the worst attainable outcome of π'_L is no worse than that of π_L. This improvement in the Leader's worst-case guarantee is not ensured by usual set-based preference relations, making the proposed formulation particularly appropriate for adversarial settings. However, alternate (weaker) notions of worst-case dominance can also be defined by adapting the relations considered in [46] to account for robustness against an adversary.

D. Irrational and Rationalizable Leader Strategies

Using the worst-case dominance relation $\overset{\text{w.c.}}{\succ}$, we define irrational and rationalizable Leader strategies as follows.

Definition 4 (Irrational Leader strategies)**.** *The set of irrational Leader strategies is:*

$$S_L^{\text{irrat}} = \left\{\pi_L \in \Pi_L \mid \exists \pi'_L \in \Pi_L \;\; \mathcal{W}_L(\pi'_L) \overset{\text{w.c.}}{\succ} \mathcal{W}_L(\pi_L)\right\}.$$

That is, a strategy π_L is deemed irrational if there exists another strategy π'_L whose worst-case outcome set dominates that of π_L according to Definition 3.

Definition 5 (Rationalizable Leader strategies)**.** *The set of rationalizable Leader strategies is the complement of the irrational set: $S_L^{\text{rat}} = \Pi_L \setminus S_L^{\text{irrat}}$.*

Equivalently, these are all Leader strategies that are *not* dominated in the worst-case sense. These rationalizable Leader strategies define a unique adversarial equilibrium concept. Unlike the established strong Stackelberg equilibrium of single-objective SSGs which makes optimistic assumptions of Follower cooperation, our framework identifies strategies that are robust and provide pessimistic, worst-case guarantees.

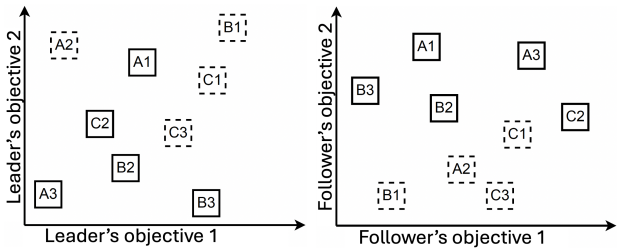


Fig. 1. Leader's (left) and Follower's (right) objective spaces in an example multiobjective game. Each square is the result of the game in which the Leader chooses a letter strategy (A, B, or C) and the Follower chooses a number strategy (1, 2, or 3). Dashed squares represent Leader-Follower strategy pairs which are not rational from the Follower's perspective, and are therefore eliminated from the Follower's optimal responses.

E. Illustration of Set-based Worst-Case Dominance

Let's consider a simple game where both the Leader and Follower have two objectives each to be maximized. The Leader must choose between three strategies (A, B, C), and the Follower must choose between three strategies (1, 2, 3). Figure 1 shows the payoffs for both players. For each Leader strategy, we determine the set of rational Follower responses: $S_F^{rat}(A) = \{1, 3\}$, $S_F^{rat}(B) = \{2, 3\}$, $S_F^{rat}(C) = \{2\}$.

Next, we identify irrational Leader strategies by comparing their worst-case outcome sets according to Definition 3. Consider strategies A and C. Strategy A yields a potential high-value outcome A1 (which dominates C's outcome C2), but also exposes the Leader to a risk of outcome A3 (if the Follower chooses response 3). Strategy C, on the other hand, effectively guarantees outcome C2. Since C2 dominates A3 (A's worst-case), we have $\mathcal{W}_L(C) \overset{\text{w.c.}}{\succ} \mathcal{W}_L(A)$, thereby providing robustness against a presumed adversarial Follower. Consequently, strategy A is deemed irrational despite its potential upside A1. Strategies B and C are mutually non-dominated, so $S_L^{rat} = \{B, C\}$.

This example demonstrates how the dominance relationships enable identification of rational strategies for both players. For simplicity, in this example we considered only pure strategies. *The same rules apply to mixed strategies as well, but instead of a player's exact payoff we compute the expected payoff with respect to the probability distribution over pure strategies.*

F. Relation to Other Multiobjective Game Formulations

Multiobjective games (MOGs) constitute a long-standing subfield grounded in vector payoffs [47], [48] and the Pareto-Nash equilibrium concept [49], [50]. While MOGs are established in the literature [51], prior methodologies rely on assumptions that our MO-SSG framework moves beyond.

Comparison to Nash Equilibria: Most MOG research rests on the simultaneous Nash equilibrium paradigm. In contrast, security domains are inherently sequential, requiring

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Transactions on Evolutionary Computation. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TEVC.2026.3749205

the Stackelberg model where the Leader commits first. This fundamental difference alters the definition of equilibria and the search methods compared to standard MOGs.

Comparison to Zero-Sum MOGs: Many approaches focus on zero-sum MOGs (e.g. [52], [53]) or mutual rationalizability [54]. Applying these to sequential SSGs is problematic, as their iterative elimination processes may prematurely discard robust strategies by failing to assume the opponent plays a payoff-optimal response. Our framework explicitly handles sequential, non-zero-sum interactions.

Another distinction is our optimization of mixed strategies, represented as a probability distribution over pure strategies. In security contexts, strategy randomization helps maintain unpredictability while maximizing one's *expected* payoff. Existing MOG formulations seldom incorporate mixed strategies as a central component of the equilibrium definition [15], [55].

Comparison to Multiobjective Bilevel Optimization: The proposed MO-SSG framework is inherently *pessimistic*, comparing Leader strategies based on the worst-case dominance relation to ensure robustness. This contrasts with the *optimistic* multiobjective bilevel optimization problem commonly studied in the classical [45] and evolutionary literature [16], [44]. Optimistic tie-breaking, as mandated by the strong Stackelberg equilibrium concept, is justified in single-objective SSGs where the Leader can induce favorable Follower responses via infinitesimal strategy adjustments [56]. However, the cooperation implied in its multiobjective variant is deemed unrealistic given the adversarial setting. Furthermore, tie-breaking in favor of the Leader allows the Leader to isolate a single most preferred response from the Follower's Pareto-optimal set, effectively ignoring the diversity of the Follower's Pareto front. As shown in [57], this possibility enables adaptive scalarization of the Follower's multidimensional payoff vector, reducing the multiobjective bilevel problem into one that is structurally a (single-objective-style) point-to-point mapping.

The positioning of our formulation relative to the existing literature on (partially) multiobjective SSGs is discussed in Section S2 of the supplementary material.

IV. Structural Properties of MO-SSGs

Before presenting our algorithmic frameworks, we establish basic properties of MO-SSGs. These foundations ensure the correctness of the evolutionary approach and provide the necessary principles for its design.

A. Evolutionary Selection under Worst-Case Dominance

The selection mechanism, central to population-based evolutionary algorithms, generally depends on the preference-based ordering of population members. Evolutionary selection pressure then iteratively biases the search towards those solutions that are considered superior under the chosen preference relation. For such an ordering to be valid in the context of MO-SSGs—i.e., to preclude cycles in the dominance structure among Leader strategies [58]—it suffices to show that the relation $\overset{\text{w.c.}}{\succ}$ satisfies *transitivity* and *irreflexivity*. Accordingly, the theorem below establishes that $\overset{\text{w.c.}}{\succ}$ indeed defines a *strict partial order*. Note that totality is not required, as evolutionary selection operates naturally on partial orders.

Theorem 1. *The worst-case dominance relation between sets of payoff vectors, $\overset{\text{w.c.}}{\succ}$, defines a strict partial order over the set Π_L of all possible Leader strategies.*

Proof. See Section S3 of the supplementary material. □

B. Sufficiency of Pure Follower Strategies

A well-established result in single-objective SSGs is that the Follower's best response to a fixed Leader strategy can always be achieved by a pure strategy [59]. This property significantly reduces the complexity of solving such games, as determining the optimum mixed strategy for the Follower would be highly computationally demanding. This restriction to pure strategies is a standard assumption widely adopted in the SSG literature (e.g. [4], [20], [27], [37] and many others), treating the Follower's decision as deterministic.

Analogously, in our MO-SSG formulation, we restrict the Follower's strategy space Π_F to the set of all pure strategies. Since the Follower makes their move as the second player, after observing the Leader's mixed strategy, they have no practical incentive for randomizing. However, because the multiobjective setting can yield multiple optimal Follower responses, the Leader faces a set of potential worst-case outcomes. We account for this non-stochastic ambiguity from a robust optimization perspective [60], as no natural probability distribution over responses is available without introducing additional assumptions. Specifically, the proposed set-based dominance relation relies on worst-case analysis to capture the uncertainty without requiring a probability distribution over Follower moves. In contrast, the Follower is able to evaluate and compare expected payoffs based on full knowledge of the Leader's committed mixed strategy. This asymmetry explains why the Follower does not require randomized moves or a worst-case (robust) formulation, whereas the Leader must anticipate a set of possible adversarial responses.

The assumed purity of Follower responses is also crucial for the computational tractability of the algorithms proposed in the next section. In the first evolutionary algorithm, it makes possible the exhaustive enumeration of S_F^{rat} for every candidate Leader strategy. In the coevolutionary algorithm, the Follower population strives to identify a representative subset of pure strategies that may reasonably approximate optimal Follower responses. Without this underlying assumption, MO-SSGs would be vastly more difficult to solve, requiring a search over the continuous space of mixed strategies for the Follower. In both algorithms described hereafter, the Leader's strategy π_L remains unrestrained (i.e., it can be either pure or mixed), as there is a strong incentive for the Leader to randomize their actions to maintain unpredictability.

V. Evolutionary Algorithms for MO-SSGs

Since a single Leader strategy maps to a *set* of optimal Follower responses rather than a unique payoff vector, standard Pareto-dominance mechanisms or decomposition-based MOEAs become ostensibly inapplicable for MO-SSGs. Algorithms must implement the worst-case dominance relation

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Transactions on Evolutionary Computation. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TEVC.2026.3749205

established in Section III to drive the selection process. To this end, we introduce two new population-based search methods for computing rationalizable Leader strategies in MO-SSGs. Both algorithms approximate the set of rationalizable Leader strategies S_L^{rat}, but differ in how the Follower's optimal responses $S_F^{\text{rat}}(\pi_L)$ are handled. The *Worst-Case Dominance Evolutionary Algorithm (WCD-EA)* assumes access to the exact Pareto set of pure Follower responses for each Leader strategy via exhaustive enumeration. The *Worst-Case Dominance Coevolutionary Algorithm (WCD-CoEA)* evolves Leader and Follower populations jointly, approximating $S_F^{\text{rat}}(\pi_L)$ online during the search. Both algorithms maintain an external archive of non-dominated Leader strategies with respect to the worst-case dominance relation $\overset{\text{w.c.}}{\succ}$.

Algorithm 1 WCD-EA for MO-SSGs.

1: Initialize Leader population $P_L \subseteq \Pi_L$ (mixed strategies)
2: **Evaluate** initial population P_L (as described below)
3: **for** generation $t = 1, 2, \ldots, T$ **do**
4: **Vary** P_L to generate offspring Q_L using binary tournament, mixture-merge crossover of mixed strategies, and mutation
5: **Evaluate** each $\pi_L \in Q_L$:
6: Compute $\mathbf{u}_L(\pi_L, \pi_F)$ and $\mathbf{u}_F(\pi_L, \pi_F)$ for all $\pi_F \in \Pi_F$
7: $S_F^{rat}(\pi_L) \leftarrow \text{ParetoPrune}\big(\{\mathbf{u}_F(\pi_L, \pi_F) : \pi_F \in \Pi_F\}\big)$
8: $\mathcal{W}_L(\pi_L) \leftarrow \text{ParetoMin}\big(\{\mathbf{u}_L(\pi_L, \pi_F) : \pi_F \in S_F^{rat}(\pi_L)\}\big)$
9: **Combine** parents and offspring: $R_L \leftarrow P_L \cup Q_L$
10: **Select** next P_L from R_L based on worst-case dominance $\overset{\text{w.c.}}{\succ}$ and set-based crowding distances
11: **end for**
12: **return** $\overset{\text{w.c.}}{\succ}$ non-dominated Leader strategies from P_L (approximation of S_L^{rat})

A. Worst-Case Dominance Evolutionary Algorithm

The first approach decouples the search over Leader and Follower strategies. For each candidate Leader strategy π_L, the Follower's optimal response set $S_F^{rat}(\pi_L)$ is computed via exhaustive evaluation of the set of all pure strategies Π_F. The Leader population is then evolved using sorting and diversity preservation operators that are specially crafted for the worst-case dominance structure of MO-SSGs. At each generation, the relative fitness of every π_L is determined by the worst-case outcome set $\mathcal{W}_L(\pi_L)$. Selection then favors strategies whose worst-case outcomes are non-dominated in the population according to the binary relation $\overset{\text{w.c.}}{\succ}$. The WCD-EA, detailed in Algorithm 1, provides a conceptually simple methodological foundation that guarantees the Follower's rationality is evaluated exactly, so that approximation errors come only from the evolutionary search on the Leader side.

Implementation Details: The implementation details of the WCD-EA for MO-SSGs are inspired by the EASG [27] method for single-objective SSGs. Each individual represents a Leader mixed strategy and the algorithm iteratively applies *crossover*, *mutation*, *evaluation*, and *selection* until a stopping condition is met. The algorithm is fully specified below.

a) Chromosome representation: Each individual encodes a Leader mixed strategy as a finite set of pure strategies with associated enactment probabilities: $\chi_q = \{(\pi_1^q, p_1^q), \ldots, (\pi_{\ell_q}^q, p_{\ell_q}^q)\}$, $\sum_{i=1}^{\ell_q} p_i^q = 1$, where ℓ_q denotes the number of pure strategies included in the mixture. In the case of security games where players interact over m time steps, a single pure strategy corresponds to a sequence of actions taken by a player across the m time steps.

b) Initial population: The initial population consists of pure strategies only ($\ell_q = 1$, $p_1^q = 1$ for all q). Each pure strategy is constructed by sampling strategies uniformly at random, independently for each chromosome.

c) Mixture-merge crossover: Recent studies in the SSG literature indicate that optimal mixed strategies for the Leader are often *sparse*, meaning a large fraction of pure strategies have a zero probability of enactment [29], [30]. Optimizing these probability distributions with standard real-coded variation operators does not enforce sparsity, slowing convergence as the operators lack mechanisms to quickly eliminate low-probability strategies. To address this, we integrate WCD-EA with a mixture–merge crossover operator which is specifically designed to maintain and promote sparsity by combining two sparse parent chromosomes into a single offspring as follows:

1) Take the union of all pure strategies from both parents and assign each probability $p_i/2$. In case of duplicates, the probabilities are summed up.
2) To maintain sparsity and prevent unbounded growth of the size of the support set (ℓ_q), remove each pure strategy (except the one with the highest probability) independently with probability $(1 - p_i)^2$.
3) Renormalize the remaining probabilities to sum to one.

This creates offspring that inherit promising strategies from both parents while maintaining a compact representation.

d) Mutation: Each offspring undergoes mutation with a fixed probability. For games played over multiple time steps, mutation selects one pure strategy inside the chromosome uniformly at random, then picks a random time step t_i and replaces all actions from t_i onward by new random actions drawn uniformly from the set of legal actions. This path–suffix rewrite ensures exploration of previously unseen trajectories while preserving part of the original strategy.

e) Evaluation: For each individual π_L, we evaluate all Follower pure strategies $\pi_F \in \Pi_F$ and compute payoff vectors $\mathbf{u}_L(\pi_L, \pi_F)$ and $\mathbf{u}_F(\pi_L, \pi_F)$. We then determine:

$$S_F^{rat}(\pi_L) = \textsc{ParetoPrune}\big(\{\mathbf{u}_F(\pi_L, \pi_F) : \pi_F \in \Pi_F\}\big),$$

$$\mathcal{W}_L(\pi_L) = \textsc{ParetoMin}\big(\{\mathbf{u}_L(\pi_L, \pi_F) : \pi_F \in S_F^{rat}(\pi_L)\}\big).$$

$\mathcal{W}_L(\pi_L)$ defines the worst–case outcome set of the Leader corresponding to strategy π_L.

Given a finite set $\mathcal{Y} \subset \mathbb{R}^{K_F}$, ParetoPrune returns the nondominated (Pareto-optimal) subset:

$$\textsc{ParetoPrune}(\mathcal{Y}) = \big\{\, \mathbf{y} \in \mathcal{Y} : \nexists\, \mathbf{y}' \in \mathcal{Y} \text{ with } \mathbf{y}' \overset{\max}{\succ} \mathbf{y} \,\big\}.$$

Given a finite set $\mathcal{Z} \subset \mathbb{R}^{K_L}$, ParetoMin returns the Pareto-minimal subset:

$$\textsc{ParetoMin}(\mathcal{Z}) = \big\{\, \mathbf{z} \in \mathcal{Z} : \nexists\, \mathbf{z}' \in \mathcal{Z} \text{ with } \mathbf{z} \overset{\max}{\succ} \mathbf{z}' \,\big\}.$$

f) Evolutionary selection: The selection operator adapts the non-dominated sorting workflow [61] to the worst–case dominance structure and point-to-set mapping of MO-SSGs:

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Transactions on Evolutionary Computation. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TEVC.2026.3749205

1) Worst-case non-dominated sorting: Given that standard Pareto dominance does not directly apply here, the partition of the population into fronts $\mathcal{F}_1, \mathcal{F}_2, \ldots$ is carried out according to the set-based worst-case dominance relation. A strategy π_L belongs to the first front $\mathcal{F}_1$ if it is not worst-case dominated by any other strategy in the combined parent and offspring population. Mathematically, $\pi_L \in \mathcal{F}_1 \iff \nexists \pi'_L \in R_L$ s.t. $\mathcal{W}_L(\pi'_L) \overset{\text{w.c.}}{\succ} \mathcal{W}_L(\pi_L)$. Once $\mathcal{F}_1$ is identified, its members are temporarily removed, and the next non-dominated front $\mathcal{F}_2$ is found among the remaining strategies. This process repeats until all individuals are assigned a rank. A valid non-dominated sorting is guaranteed by the strict partial order defined by $\overset{\text{w.c.}}{\succ}$ according to Theorem 1.
2) The next generation of individuals is selected front by front. If the inclusion of all individuals in a particular front causes the population limit to be exceeded, a niching measure—computed using a single point representative of the set $\mathcal{W}_L(\pi_L)$ (see discussion below on *set-based crowding distances*)—is employed to select only a subset of individuals from that front.

These steps give rise to elitist selection and diversification in the Leader's objective space, facilitating convergence toward a diverse set of rationalizable Leader strategies. The final non-dominated set of strategies is then returned as the approximation to S_L^{rat}. The dominant cost per generation of the WCD-EA comes from the exhaustive enumeration of optimal Follower responses, but skyline pruning ensures only minimal points are stored in $\mathcal{W}_L(\pi_L)$. The number of extreme points is typically small in practice, keeping evaluation computationally tractable.

g) Set-based crowding distances: Niching measures for diversity maintenance are usually defined among individual vectors, whereas MO-SSGs give rise to sets of payoff vectors. In order to enable diversity preservation under point-to-set mappings, we represent each Leader strategy π_L by the *nadir point* of its worst-case outcome set $\mathcal{W}_L(\pi_L)$. Formally, if $\mathcal{W}_L(\pi_L) = \{\mathbf{u}^{(1)}, \mathbf{u}^{(2)}, \ldots\} \subset \mathbb{R}^{K_L}$, then the nadir vector $\mathbf{v}^{rep}(\pi_L) \in \mathbb{R}^{K_L}$ is computed as: $v_k^{rep}(\pi_L) = \min_{j=1,2,\ldots}\{u_k^{(j)}\}, \ \forall k \in \{1, \ldots, K_L\}$. The theoretical rationale behind using the nadir point as a proxy for niching is grounded in the observation that worst-case dominance between sets relates to weak Pareto dominance between their respective nadir points. (No such monotonic relationship between worst-case dominance and Pareto dominance holds for the ideal point.) We formalize this observation through a theorem adapted from [14].

Theorem 2. *Let $\mathbf{v}^{rep}(\pi_L)$ denote the nadir point of the worst-case outcome set $\mathcal{W}_L(\pi_L)$. If a Leader strategy π_A worst-case dominates strategy π_B (i.e., $\mathcal{W}_L(\pi_A) \overset{\text{w.c.}}{\succ} \mathcal{W}_L(\pi_B)$), then $\mathbf{v}^{rep}(\pi_A)$ Pareto-dominates ($\overset{max}{\succ}$) or is equal to $\mathbf{v}^{rep}(\pi_B)$. In short, $\mathbf{v}^{rep}(\pi_A)$ weakly Pareto-dominates $\mathbf{v}^{rep}(\pi_B)$.*

Proof. See Section S3 of the supplementary material. □

With a representative point identified for each set $\mathcal{W}_L(\pi_L)$, it becomes possible for off-the-shelf niching measures to be used to quantify strategy diversity in the Leader's objective space. In our implementation, we adopt the *crowding distance* from [61], preferentially selecting individuals with higher values to maintain a broad distribution of strategies.

B. Worst-Case Dominance Coevolutionary Algorithm

The second approach couples the search dynamics of Leader and Follower strategies in a *competitive coevolutionary* process. A pseudo-code is provided in Algorithm 2, in which two populations, P_L and P_F, are evolved simultaneously:

- The Leader population P_L seeks strategies with minimal worst-case losses against the current Follower population.
- The Follower population P_F seeks a subset of pure strategies that may reasonably approximate optimal responses.

Algorithm 2 WCD-CoEA for MO-SSGs.

```
Initialize populations P_L (Leader, mixed) and P_F (Follower, pure)
Evaluate initial populations against each other
for epoch e = 1, 2, ..., E do
    Follower phase (evolve against frozen P_L):
        Generate offspring Q_F from P_F
        Evaluate Q_F against all π_L ∈ P_L
        Combine R_F ← P_F ∪ Q_F and Select next P_F (using Pareto dominance and crowding distances for niching)
    Leader phase (evolve against frozen P_F):
        Generate offspring Q_L from P_L
        Evaluate Q_L: compute W_L(π_L) against all π_F ∈ P_F
        Combine R_L ← P_L ∪ Q_L and Select next P_L (using worst-case dominance ≻^w.c. and set-based crowding distances)
end for
return ≻^w.c. non-dominated Leader strategies from P_L
```

At each generation, every $\pi_L \in P_L$ is evaluated against all $\pi_F \in P_F$ to approximate its worst-case outcome set. Simultaneously, every $\pi_F \in P_F$ is evaluated against all $\pi_L \in P_L$, and selection within P_F favors strategies that are performant on average against the evolving set of Leader strategies. Over time, P_F tends to converge towards generalizable responses, while P_L converges to strategies with non-dominated worst-case payoffs conditioned on P_F. This coevolutionary process approximates the two-player game dynamics without full enumeration of Π_F, typically reaching high-quality approximations of S_L^{rat} with fewer overall payoff evaluations.

Implementation Details: The coevolutionary algorithm maintains two interacting populations: one for the Leader's mixed strategies and one for the Follower's pure strategies. Both populations evolve alternately, influencing each other's fitness landscapes. Below we describe only those components of the WCD-CoEA that differ from the single-population WCD-EA described in the previous subsection.

The Follower population serves two purposes: (i) approximating the Leader's worst-case outcome set by providing candidate responses, and (ii) evolving over time to exploit weaknesses in the Leader population, thus guiding the Leader to correct their weaknesses in the direction of rationalizable strategies. *This competitive interaction reduces the need for exhaustive Follower enumeration at each generation.*

a) Alternating evolution: The algorithm proceeds in *epochs*. In each epoch, we first evolve the Follower population for a fixed number of generations while keeping the Leader

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Transactions on Evolutionary Computation. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TEVC.2026.3749205

population fixed. Then the Leader population is evolved for the same number of generations against the updated Follower population. This alternating scheme gradually strengthens both sides: the Follower set converges towards generalizable responses, and the Leader set learns to be robust against the vulnerabilities they expose.

b) Crossover and mutation in the Follower population: Since Follower individuals are pure strategies, their variation operators differ from the Leader side. Consider games where players interact over m time steps, such that a pure strategy is a sequence of actions taken across the m time steps. Then, *crossover* in the Follower population is one-point crossover and exchanges suffixes of two parent strategies after the first common action in time. If no common action exists, parents remain unchanged. In *mutation*, starting from a randomly selected time step, all subsequent actions are replaced by uniformly drawn legal actions, introducing novel trajectories.

c) Evaluation: For each Leader individual π_L, its worst-case multiobjective payoff set $\mathcal{W}_L(\pi_L)$ is approximated against all Follower individuals in the current population P_F. As the Follower population evolves, this approximation may move closer to the true worst-case set. Conversely, each Follower individual π_F is evaluated against the entire current Leader population P_L. Unlike WCD-EA, a single Follower is pitted against multiple Leader strategies. We therefore employ a simplified fitness measure given by the average payoff vector: $\bar{\mathbf{u}}_F(\pi_F) = \frac{1}{|P_L|}\sum_{\pi_L \in P_L} \mathbf{u}_F(\pi_L, \pi_F)$, which quantifies the generalizability of a response. Although this measure does not strictly conform with the definition of optimal Follower responses, the resulting WCD-CoEA helps achieve computational speedups for time-sensitive security applications that would be difficult to achieve with exhaustive enumeration. The averaged payoff can be regarded as a low-fidelity "surrogate" objective function that significantly reduces evaluation overhead. However, the averaging does introduce a potential bias: the Leader may underestimate the true severity of the Follower's worst-case response, which can occasionally lead to convergence upon Leader strategies that are sub-optimal in a strict worst-case sense.

d) Selection: Both Leader and Follower populations undergo non-dominated sorting and niching-based environmental selection similar to what was considered in the WCD-EA. The Leader population is sorted by the worst-case dominance relation $\overset{\text{w.c.}}{\succ}$ determined according to the current approximation to $\mathcal{W}_L(\pi_L)$. The selection operator acting on the Follower population makes use of Pareto dominance of the averaged payoff vector $\bar{\mathbf{u}}_F$, favoring Followers that perform well on average against the current spectrum of Leader strategies.

The final non-dominated Leader strategies are returned as an approximation of the rationalizable set S_L^{rat}.

C. Query Complexity Analysis

The computational cost per generation (or epoch) of the WCD-EA and WCD-CoEA is dominated by their query complexity for strategy evaluations. Let N be the population size (the number of offspring generated per generation) and let K_L and K_F be the number of objectives for the Leader and Follower, respectively. Consider each objective to constitute an independent function evaluation. In the WCD-EA, obtaining $\mathcal{W}_L(\pi_L)$ for a single Leader strategy requires exhaustive evaluation of the entire pure strategy space, Π_F, of the Follower. Since both the Leader's and Follower's payoffs must be calculated to identify $\mathcal{W}_L(\pi_L)$, a single evolutionary step of the Leader in WCD-EA performs $(K_L + K_F)\ N\ |\Pi_F|$ function evaluations. In contrast, the WCD-CoEA evaluates the Leader's strategy exclusively against the current Follower population, P_F (where $|P_F| = N$). Consequently, an equivalent evolutionary step in WCD-CoEA requires $(K_L + K_F)\ N^2$ function evaluations. By setting $N << |\Pi_F|$, computational advantages can therefore be gained for the coevolutionary algorithm but at the cost of reduced precision in the computed Follower responses.

Furthermore, the cost of a single function evaluation, $u(\pi_L, \pi_F)$, is influenced by the specific game representation, its parameters, and the size of the support set of the mixed strategy π_L. A mixed strategy is represented as a probability distribution over l_q pure strategies. To compute the expected payoff, the algorithm must simulate the interaction for each pure strategy in this mixture against the given Follower strategy. Each simulation is an independent query to a black-box or lookup table. Since l_q can be no greater than $|\Pi_L|$, the overall query complexity of the WCD-EA and the WCD-CoEA in the worst case becomes $(K_L + K_F) \cdot N \cdot |\Pi_F| \cdot |\Pi_L|$ and $(K_L + K_F) \cdot N^2 \cdot |\Pi_L|$, respectively.

In games where a pure strategy is characterized by a sequence of actions taken over m time steps, the cost of a single query may scale as $O(m)$. The dominant computational load in each generation of the two algorithms then scales as $O((K_L + K_F) \cdot N \cdot |\Pi_F| \cdot |\Pi_L| \cdot m)$ for the WCD-EA and $O((K_L + K_F) \cdot N^2 \cdot |\Pi_L| \cdot m)$ for the WCD-CoEA.

D. Adapting Standard MOEAs as Baselines

To comprehensively assess the performance of WCD-EA and WCD-CoEA, well-established MOEAs, NSGA-III [62] and MOEA/D [63], are adapted to serve as comparative baselines. Recall that the MO-SSG formulation yields a point-to-set mapping that precludes direct applicability of Pareto-based MOEAs. To bridge this gap, we implement two distinct baseline adaptations. The first adaptation leverages the insight from Theorem 2. The point-to-set mapping of a Leader strategy π_L is reduced to the nadir point $\mathbf{v}^{rep}(\pi_L)$ of its worst-case outcome set $\mathcal{W}_L(\pi_L)$. This single multi-dimensional point serves as the fitness value for the given strategy, allowing standard Pareto dominance to be used without further modification. The second adaptation employs standard set-based dominance (Definition 2) directly to compare the worst-case outcome sets. This allows the evolutionary selection pressure to operate on sets rather than on single representative points, albeit without the strict pessimistic guarantees of the proposed worst-case dominance relation.

NSGA-III and MOEA/D have been adapted from the `pymoo` Python library [64]. It is worth highlighting that while worst-case dominance implies weak Pareto dominance of nadir points (Theorem 2), vice versa is not necessarily true. Also, recall that worst-case dominance between sets is

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Transactions on Evolutionary Computation. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TEVC.2026.3749205

generally, but with important exceptions, logically independent of standard set-based dominance. As such, the evolutionary selection pressure induced in the baseline algorithms may not precisely match that of the robust WCD-EA. The implications of this mismatch are examined in the experimental results.

VI. Experimental Setup

We consider two distinct benchmark domains to assess the performance of the proposed algorithms. The first is an instructional *Tug of War* (ToW) game whose geometric nature allows the analytical derivation and visualization of optimal solutions, enabling a rigorous evaluation of the algorithms against the exact ground truth. Our analysis also shows that this game yields optimal mixed strategies that are highly sparse, which is a prevalent feature in real-world SSGs. It therefore serves as a test for an algorithm's ability to discover and maintain solution sparsity. The second example extends the well-established *FlipIt* security game [29] to a multiobjective setting. This represents a highly combinatorial strategy space typical of cybersecurity domains. Operating on graph-based network topologies under incomplete information, it tests the algorithms' scalability and robustness in resolving multifaceted trade-offs between long-term control and limited resources.

A. The Tug of War Game

We adapt the multiobjective ToW game proposed in [53] by accounting for mixed strategies by the Leader. This game models a scenario where a Leader (e.g., military forces, park rangers) defends a high-value central asset (e.g., power plant, water source) against a Follower (e.g., terrorists, poachers) attempting to breach the perimeter. The interaction in ToW is modeled abstractly as a contest of applied forces in a multi-dimensional space, with the outcome being a net displacement vector representing the asset's control. Similar ToW inspired games have been widely used for evaluation in other multiobjective optimization studies, e.g., [14], [52], [55].

1) Basic Game Description: We initially present the game mechanics in a two-dimensional setting. This formulation serves as a baseline and is naturally generalized to higher-dimensional spaces. In this base case, the interaction is modeled as a contest of forces in a 2D plane.

- **Leader:** Chooses an angle $\theta_L \in [0, 360^\circ)$ and a force magnitude $f_L \in [0, 1]$. This gives a force vector $\vec{F}_L = (f_L \cos\theta_L, f_L \sin\theta_L)$.
- **Follower:** Chooses $\theta_F \in [0, 360^\circ)$ and $f_F \in [0, 1]$, giving $\vec{F}_F = (f_F \cos\theta_F, f_F \sin\theta_F)$.
- **Outcome:** Assuming the forces to be applied for a duration of $\sqrt{2}$ units of time, the result is a displacement vector $\vec{R} \in \mathbb{R}^2$ that represents the net control of the asset:

$$\vec{R} = \vec{F}_L - \vec{F}_F. \tag{1}$$

The Leader aims to maximize the displacement components R_x and R_y (i.e., $\mathbf{u}_L = \vec{R}$), while the Follower seeks to minimize them (i.e., maximize the negation of $\vec{R}$ such that $\mathbf{u}_F = -\vec{R}$), forming a zero-sum game.

To visualize these dynamics, Figure 2 depicts a sample interaction between players. The Leader selects a strategy defined by the angle θ_L and magnitude corresponding to the vector $\vec{F}_L$, while the Follower exerts force $\vec{F}_F$ in direction θ_F. The final game outcome $\vec{R} = (R_x, R_y)$ represents the net displacement resulting from these competing forces assuming they are applied at the origin $(0, 0)$. The Leader effectively tries to pull the outcome into the first quadrant towards the perimeter of the unit circle (denoted by the dotted arc).

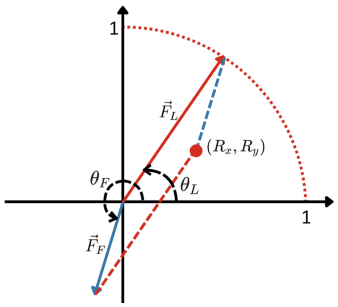


Fig. 2. Geometric representation of the Tug of War game mechanics. The outcome (R_x, R_y) (red dot) is determined by the vector interaction between the Leader's force $\vec{F}_L$ (red arrow) and the Follower's force $\vec{F}_F$ (blue arrow).

2) Problem Discretization: Since the vast majority of existing Security Games operate in discrete action spaces, we chose to discretize the strategy space of the ToW game. We define three levels of discretization: *Coarse*, *Medium*, and *Fine*, which allow us to study the algorithms' performance as the problem scale increases. The discretization determines the number of available angles (N_θ) and force levels (N_f) for both players. The total number of pure strategies available to each player is therefore $N_\theta \times N_f$. The configurations are defined as: **Coarse:** 12 angles (30° step), 5 force levels (0.2 step). Total pure strategies: 60. **Medium:** 36 angles (10° step), 10 force levels (0.1 step). Total pure strategies: 360. **Fine:** 72 angles (5° step), 20 force levels (0.05 step). Total pure strategies: 1440.

3) Dimensionality of Vector Payoffs: We also define two different versions of the problem based on the number of objectives ($K \in \{2, 3\}$), representing the dimensionality of the space in which the ToW game is played.

a) Bi-objective case ($K = 2$), described above: The Leader aims to maximize control along two orthogonal axes:

$$\max \mathbf{u}_L = (R_x, R_y).$$

The Follower aims to minimize these values, i.e., $\max \mathbf{u}_F = -(R_x, R_y)$, as shown by the interaction in Figure 2.

b) Tri-objective case ($K = 3$): We extend the interaction to a three-dimensional space, where the Leader aims to maximize control along three orthogonal axes:

$$\max \mathbf{u}_L = (R_x, R_y, R_z).$$

In this setting, players define their strategies using spherical coordinates: force magnitude $f \in [0, 1]$, azimuth angle $\theta \in [0, 360^\circ)$, and polar angle $\phi \in [0, 180^\circ]$. Consequently, the force vector components are calculated as $\vec{F} = (f \sin\phi \cos\theta, f \sin\phi \sin\theta, f \cos\phi)$, and the outcome $\vec{R}$ is the vector difference of these forces in 3D space.

4) Characterizing Optimal Mixed Strategies of the Leader: From our analysis of the 2D ToW game in Section S4 of the

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Transactions on Evolutionary Computation. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TEVC.2026.3749205

supplementary material, it follows that the set of **rationalizable Leader strategies consists exclusively of pure strategies or mixtures of only two adjacent strategies** (with respect to the choice of angle in $[0, 90°]$), **always employing the maximum force** $f_L = 1$. In the discretized game variant, the mixed strategies play an important role in filling the gaps between the discrete angular sectors. Discovering these *highly sparse* mixtures via an evolutionary algorithm is deemed non-trivial. A similar analysis revealing the sparsity of optimal mixed strategies applies to the 3D ToW game setting as well. There, the optimal mixtures arise from randomization over only those neighboring angles whose outcome vectors (for any pure Follower strategy) lie on a common triangular face of the convex hull's extreme surface.

A distinctive property of the structure of the Leader's worst-case outcome set in the ToW game is that worst-case dominance, standard set-based dominance, and nadir point-based dominance logically coincide. Consequently, ToW provides a controlled benchmark in which the effectiveness of the sparsity-preserving (mixture-merge) variation operator can be assessed without conflating its effect with differences between these dominance relations.

B. The FlipIt Security Game

Next, we consider security scenarios represented by FlipIt Games (FIGs) [65]. FIGs model adversarial interactions where two players repeatedly compete for control over networked resources. Inspired by cybersecurity, the *Follower* seeks to compromise system components such as servers or routers, while the *Leader* aims to detect these compromised components and reclaim control. FIGs capture the stealthy and ongoing nature of cyber conflicts, where control may change hands without immediate detection.

A FIG is played on a directed graph with n vertices for m discrete time steps. Each vertex represents a controllable system component, and edges define possible propagation paths. At each time step, both players *simultaneously* select one vertex to attempt to take control of—a move called a *flip*. Initially, all vertices are under the Leader's control, while a subset of *entry nodes*, representing exposed network interfaces, is directly accessible to the Follower. A flip attempt succeeds if (1) the player already controls at least one predecessor of the chosen node (unless it is an entry node), and (2) the opponent does not attempt to flip the same node in that step.

Each vertex v_i provides a reward $\mathbf{r}_i > 0$ for maintaining control and incurs a cost $\mathbf{c}_i > 0$ for performing a flip. The cumulative *payoff* of each player is the total reward from controlled nodes minus the total cost of all flip attempts, whether successful or not. Players must balance the long-term control with limited resources, leading to dynamic trade-offs between aggressive and conservative strategies.

In this work, we employ the *No-Info* variant of FIGs, in which players act without knowledge of the opponent's current control state or recent actions. This setting captures realistic cybersecurity situations with incomplete or delayed information. The No-Info FIG benchmark has been widely adopted in previous security game studies [28], [36] and we follow this convention. An illustration of a FIG scenario is provided in Section S5 of the supplementary material.

We take the FIG benchmark graphs from earlier works and keep their structural parameters (namely graph topologies, time horizons, and entry-node sets) unchanged. This means 5 instances for each (n, m) with $n \in \{5, 10, 15, 20, 25, 30, 40\}$ and $m \in \{3, 4, 5, 6, 8, 10, 15, 20\}$, yielding 280 game instances in total. All game instances were grouped by size according to the total problem scale $(n \times m)$: small for $(n, m) \in \{(5-10, 3-5)\}$, medium for $(n, m) \in \{(15-25, 6-10)\}$, and large for $(n, m) \in \{(30-40, 15-20)\}$. To study the multiobjective setting, for each node we randomly generate a $(K-1)$-dimensional payoff vector (reward components for control and cost components for flip attempts) and add the original payoff as the K-th element of the vector. Thus, we increase the number of objectives while leaving the graphs unchanged. We consider $K \in \{2, 4, 8\}$ in all tests, and the Leader and Follower have the same number of objectives.

C. Evaluation Protocol and Metrics

Each game instance is run for 20 independent trials per method with distinct seeds. For all the algorithms included in the experimental study, we set the stopping criterion to a fixed number of Leader generations equal to 200. In the case of WCD-CoEA, which proceeds in epochs of alternating evolution (with 10 Leader generations and 10 Follower generations per phase), this corresponds to running the algorithm until the Leader population has undergone 200 evolutionary steps in total. Leader strategies produced by WCD-CoEA are re-evaluated using exact exhaustive enumeration of Follower strategies at the end of the run.

We did not perform extensive parameterization, as result optimization was not the main focus of this paper. Most parameter values were taken from recommended values in single-objective SSG papers [27], [28]. Namely, population size $N = 200$ and crossover probability $p_c = 0.9$.

We employ the *Inverted Generational Distance* (IGD) as our primary metric for all results and comparisons. IGD measures the convergence (in objective space) of the evolved strategy set to a reference set that represents the ground truth, simultaneously assessing both proximity and diversity. For the ToW benchmark, the rationalizable strategies of the Leader (and the corresponding optimal Follower responses) have been analytically determined, allowing us to compute IGD directly against the known rational game outcomes. In the case of FlipIt games, where the analytical optima are unavailable, we attempt to approximate a representative set of rational game outcomes. To ensure an unbiased evaluation, this reference set is constructed by pooling the final populations from extended executions (running for 5000 generations) of all evaluated algorithms. We then extract the non-dominated solutions from this combined pool with respect to the worst-case dominance relation. We report results averaged over 20 independent trials for each game instance. All experiments were run on Intel Xeon Silver 4116 @ 2.10 GHz with 256 GB RAM. Source code and generated FIG instances (graphs, rewards/costs) are released at https://github.com/zychowskia/MultiobjectiveSSG.

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Transactions on Evolutionary Computation. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TEVC.2026.3749205

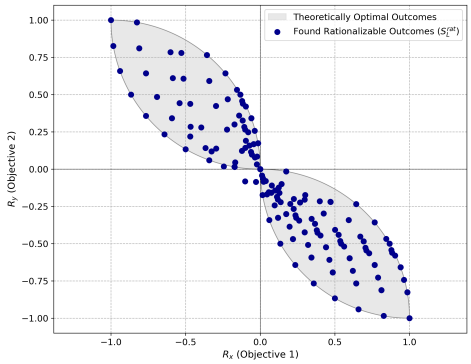


Fig. 3. Visualization of the worst-case outcomes of the rationalizable Leader strategies found by WCD-EA for the 2D ToW game.

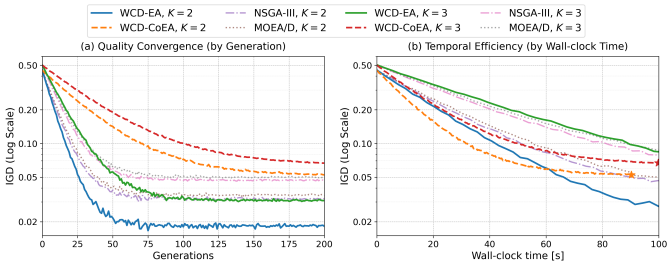


Fig. 4. IGD convergence analysis (log scale) of WCD-EA, WCD-CoEA, and the standard MOEA baselines (NSGA-III, MOEA/D) for the ToW game ($K = 2, 3$), averaged across all discretization levels. Left subplot displays quality convergence over generations. Right subplot illustrates performance with respect to wall-clock time (with the horizontal axis truncated to 100 seconds).

VII. Experimental Results

This section compares the WCD-EA (with exhaustive Follower response evaluation) and the WCD-CoEA (with a co-evolved Follower population) in terms of quality and computation time. Results are averaged over all runs and games.

A. Tug of War Games

We study WCD-EA and WCD-CoEA on ToW, focusing on their ability to approximate the set of analytically determined optimal game outcomes (see Definition 6 below for the 2D game version) and the rationalizable Leader strategies described in Section VI-A4. Table I summarizes the results.

Definition 6 (Optimal ToW Game Outcomes)**.** *Let $\mathcal{L}$ and $\mathcal{F}$ denote the sets of strategies for the Leader and the Follower, respectively, represented as arcs of the unit circle in the first and third quadrants:*

$$\mathcal{L} = \{(\cos\theta_L, \sin\theta_L) \in \mathbb{R}^2 \mid 0 \leq \theta_L \leq \frac{\pi}{2}\} \tag{2}$$

$$\mathcal{F} = \{(\cos\theta_F, \sin\theta_F) \in \mathbb{R}^2 \mid \pi \leq \theta_F \leq \frac{3\pi}{2}\} \tag{3}$$

The region of all possible game outcomes $\mathcal{R}$ is the Minkowski sum of these sets:

$$\mathcal{R} = \mathcal{L} \oplus \mathcal{F} = \{\mathbf{l} + \mathbf{f} \mid \mathbf{l} \in \mathcal{L}, \mathbf{f} \in \mathcal{F}\} \tag{4}$$

The boundary of the region $\mathcal{R}$ is defined by four circular arcs of radius $r = 1$ centered at points $(\pm 1, 0)$ and $(0, \pm 1)$.

As shown in Table I, the WCD-EA achieves a lower IGD to the ground truth (Definition 6) than all other methods, with Welch's t-test confirming these differences to be statistically significant ($p < 0.01$) across all ToW configurations. All the adapted baseline MOEAs in this experiment are implemented using standard, off-the-shelf variation operators from `pymoo`. This serves to isolate the impact of sparsity-preserving variation, as the dominance relations used by all algorithms logically coincide for ToW. The resulting performance gap highlights the importance of our specialized mixture-merge mechanism. For completeness, the supplementary material includes the results of the baseline MOEAs equipped with the same mixture-merge crossover, confirming their near-identical convergence when specialized mechanisms are applied.

Between the two bespoke methods, WCD-CoEA achieves an IGD level that's inferior to that of WCD-EA, but offers substantial speedups especially for medium and fine configuration games, in line with discussions in Section V-C. For the special case of coarse configurations, the temporal gain is seen to reverse only because the population size N exceeds $|\Pi_F|$.

Figure 3 visualizes the objective space for a representative run of the WCD-EA on the bi-objective instance. The set of rationalizable outcomes identified by the WCD-EA (represented by blue points) are well distributed within the boundary of the theoretically optimal region from Definition 6 (shown in gray). This empirical result confirms that the algorithm successfully converges to the theoretical optimum.

1) Convergence Rate Analysis: To better understand the optimization dynamics, we analyze the convergence of the IGD metric. Figure 4 illustrates the IGD trajectories of all tested algorithms for both bi-objective ($K = 2$) and tri-objective ($K = 3$) formulations. The results in Figure 4(a) confirm that WCD-EA (solid lines) consistently achieves superior convergence in terms of generations, dropping quickly to IGD values below 0.05. NSGA-III and MOEA/D experience a comparably steep initial IGD drop, but plateau early at suboptimal levels. This premature stagnation is primarily associated with their variation operators that are not designed to preserve sparsity. The superiority of WCD-EA over WCD-CoEA indicates that exact Follower strategy enumeration provides a consistent selection pressure for the Leader to utilize for optimization. Interestingly, while WCD-CoEA (dashed lines) may seem to converge slower by generations, its potential computational gains over all other methods in terms of wall-clock time is evident in Figure 4(b), especially in the tri-objective example.

2) Analysis of Rationalizable Leader Strategies: We further examine the structure of the solutions found by the WCD-EA to verify if they are in agreement with the analytical prediction, i.e., optimal mixed strategies must employ maximum force in the first quadrant and, in the 2D case, be composed of at most two adjacent pure strategies.

Figure 5 presents a polar heatmap of the Leader strategies accumulated over the final population. The heatmap reveals a high concentration of strategies (indicated by bright yellow color) strictly located on the outer rim, corresponding to a force magnitude of 1. Furthermore, these strategies are

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Transactions on Evolutionary Computation. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TEVC.2026.3749205

TABLE I

COMPARISON OF WCD-EA, WCD-COEA, NSGA-III, AND MOEA/D ON TOW ($K = 2, 3$) IN TERMS OF QUALITY (IGD - LOWER IS BETTER) AND COMPUTATION TIME (SECONDS - LOWER IS BETTER). TO HIGHLIGHT THE NECESSITY OF SPARSITY PRESERVING MECHANISMS, THE BASELINE MOEAS (NSGA-III AND MOEA/D) ARE EVALUATED HERE USING THE STANDARD, OFF-THE-SHELF VARIATION OPERATORS IN PYMOO. MEAN ± STD DEV ARE REPORTED WITH THE BEST RESULTS **BOLDED**.

		IGD				Time (s)			
K	Level	WCD-EA	WCD-CoEA	NSGA-III	MOEA/D	WCD-EA	WCD-CoEA	NSGA-III	MOEA/D
2	Coarse	**0.012** $\pm$ 0.002	0.035 $\pm$ 0.005	0.022 $\pm$ 0.003	0.024 $\pm$ 0.004	**3.2** $\pm$ 0.2	10.5 $\pm$ 0.5	3.4 $\pm$ 0.3	3.3 $\pm$ 0.2
	Medium	**0.018** $\pm$ 0.003	0.052 $\pm$ 0.008	0.031 $\pm$ 0.005	0.034 $\pm$ 0.006	117.5 $\pm$ 7.4	**65.3** $\pm$ 4.1	119.2 $\pm$ 7.8	118.6 $\pm$ 7.5
	Fine	**0.025** $\pm$ 0.004	0.068 $\pm$ 0.010	0.043 $\pm$ 0.006	0.046 $\pm$ 0.007	1405.4 $\pm$ 90.0	**195.2** $\pm$ 12.5	1412.8 $\pm$ 91.5	1409.1 $\pm$ 89.2
3	Coarse	**0.020** $\pm$ 0.003	0.045 $\pm$ 0.006	0.032 $\pm$ 0.004	0.035 $\pm$ 0.005	**3.5** $\pm$ 0.2	11.8 $\pm$ 0.6	3.7 $\pm$ 0.3	3.6 $\pm$ 0.2
	Medium	**0.032** $\pm$ 0.005	0.062 $\pm$ 0.009	0.047 $\pm$ 0.007	0.050 $\pm$ 0.008	129.8 $\pm$ 9.5	**72.1** $\pm$ 5.3	132.4 $\pm$ 9.8	130.7 $\pm$ 9.1
	Fine	**0.041** $\pm$ 0.006	0.085 $\pm$ 0.012	0.062 $\pm$ 0.008	0.065 $\pm$ 0.010	1550.9 $\pm$ 106.5	**215.4** $\pm$ 14.8	1564.2 $\pm$ 108.1	1555.3 $\pm$ 105.8

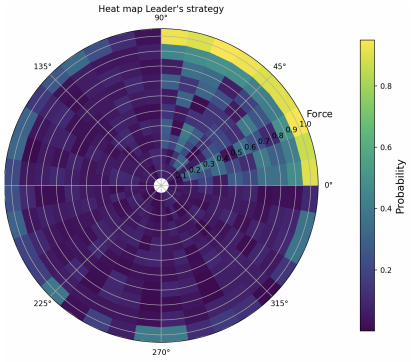


Fig. 5. Polar heatmap of the evolved Leader strategies by WCD-EA for the Medium ToW game in 2D. The radial axis represents force magnitude (0.0 to 1.0), and the angular axis represents the chosen sector. The high concentration of probability mass (yellow regions) on the outer rim between 0° and 90° confirms that the evolved rationalizable strategies are those that employ maximal force in the first quadrant.

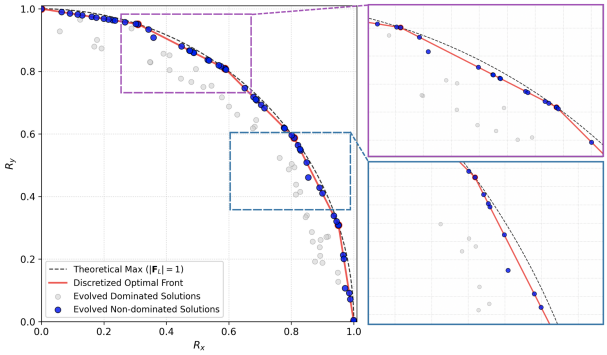


Fig. 6. Visualization of evolved Leader strategies by WCD-EA in the bi-objective ToW game ($K = 2$). The blue points represent the outcomes of the Leader non-dominated strategies assuming no Follower resistance. They spread along the red polygonal chain, which represents the optimal frontier attainable under strategy discretization, slightly below the theoretical continuous maximum (dashed black arc).

distributed almost exclusively in the first quadrant, between $0°$ and $90°$. This distribution reinforces that strategies with partial force ($f_L < 1$) are dominated, and the Leader (optimally) distributes resources only in the directions that positively contribute to the objectives R_x and R_y, ignoring sectors that would result in negative objective components.

We have also quantified the sparsity of the evolved solutions by analyzing their average support size (i.e., the number of pure strategies with non-zero probability in a mixture). The observed average support of 1.86 (out of a total of 60, 360 and 1440 available pure strategies in the coarse, medium and fine configurations, respectively) is consistent with our theoretical analysis that rationalizable Leader strategies in the 2D ToW game consist exclusively of pure strategies and highly sparse mixtures of only two adjacent strategies. The ability of the WCD-EA to achieve such sparsity is deemed a direct consequence of the mixture-merge crossover operator.

To investigate the structural properties of the solutions found, we visualize the evolved Leader strategies in the objective space. Figure 6 presents the distribution of Leader strategies discovered by the WCD-EA for a bi-objective instance ($K = 2$). The visualization assumes a simplified interaction scenario where the Follower exerts zero force ($\vec{F}_F = 0$), enabling a direct mapping of the Leader's strategy vector $\vec{F}_L$ to the game outcome. The theoretical frontier in a *continuous* strategy space would correspond to the arc of the unit circle in the first quadrant (dashed black curve). However, due to the discretization of the action space, the attainable efficient frontier is formed by the maximizing edge of the convex hull of discrete pure strategies, represented by the chords connecting adjacent optimal points (red polygonal chain). The blue points depicting the evolved population map tightly to our theoretical deduction, effectively mixing adjacent pure strategies and maximizing force magnitude within the limitations of the discrete domain. Similar results are also obtained for the 3D game, as shown in Section S6 of the supplementary material.

B. FlipIt Games

In what follows, we examine algorithmic performance on FlipIt games grounded in the domain of cybersecurity.

1) Quality: Table II summarizes the results for all instance sizes and objective counts K. WCD-EA consistently achieves lower IGD values across all settings, confirming its advantage in approximating rationalizable Leader strategies. However, this accuracy comes at a higher computational cost compared to the WCD-CoEA. Both methods exhibit increasing IGD and increasing runtime with larger instance sizes and higher K, which illustrates the growing complexity with higher objective space dimensionality in the extended FlipIt game settings.

Unlike ToW, the complex combinatorial landscape of the FlipIt game means that the underlying dominance relations may no longer coincide. To provide a robust evaluation, we include two baseline algorithm variants, one utilizing standard set-based dominance (Definition 2) and the other using nadir-

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Transactions on Evolutionary Computation. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TEVC.2026.3749205

TABLE II

COMPARISON OF WCD-EA, WCD-CoEA, NSGA-III, AND MOEA/D ON FLIPIT (MEAN ± STD DEV) IN TERMS OF QUALITY (IGD - LOWER IS BETTER) AND COMPUTATION TIME (SECONDS - LOWER IS BETTER). BASELINES ARE PRESENTED IN TWO VARIANTS: NADIR-POINT APPROXIMATIONS AND STANDARD SET-BASED DOMINANCE (DEFINITION 2). ALL BASELINE METHODS UTILIZED THE SPECIALIZED MIXTURE-MERGE CROSSOVER. THE BEST RESULTS IN EACH CATEGORY ARE **BOLDED**.

		IGD						**Time (s)**					
		Proposed (Worst-Case)		**Baseline (Nadir)**		**Baseline (Set-based)**		**Proposed (Worst-Case)**		**Baseline (Nadir)**		**Baseline (Set-based)**	
K	**Group**	**WCD-EA**	**WCD-CoEA**	**NSGA-III**	**MOEA/D**	**NSGA-III**	**MOEA/D**	**WCD-EA**	**WCD-CoEA**	**NSGA-III**	**MOEA/D**	**NSGA-III**	**MOEA/D**
2	Small	**0.012** ± 0.003	0.036 ± 0.006	0.045 ± 0.006	0.049 ± 0.007	0.038 ± 0.005	0.041 ± 0.006	22.6 ± 1.2	**8.7** ± 0.3	23.1 ± 1.2	22.8 ± 1.1	23.4 ± 1.2	23.1 ± 1.1
	Medium	**0.014** ± 0.003	0.043 ± 0.007	0.054 ± 0.007	0.058 ± 0.008	0.046 ± 0.007	0.049 ± 0.007	39.1 ± 1.9	**11.4** ± 0.5	39.4 ± 2.0	38.9 ± 1.9	39.8 ± 2.1	39.2 ± 1.9
	Large	**0.017** ± 0.004	0.051 ± 0.009	0.063 ± 0.009	0.068 ± 0.010	0.055 ± 0.008	0.058 ± 0.009	61.2 ± 2.9	**15.6** ± 0.6	62.1 ± 2.9	61.2 ± 2.8	62.8 ± 3.0	61.9 ± 2.9
4	Small	**0.039** ± 0.005	0.108 ± 0.013	0.122 ± 0.014	0.128 ± 0.015	0.108 ± 0.013	0.115 ± 0.014	51.9 ± 2.5	**12.4** ± 0.6	52.6 ± 2.4	51.8 ± 2.3	53.3 ± 2.5	52.5 ± 2.4
	Medium	**0.044** ± 0.006	0.115 ± 0.016	0.131 ± 0.015	0.138 ± 0.016	0.119 ± 0.014	0.125 ± 0.015	82.3 ± 4.1	**17.1** ± 0.8	83.5 ± 4.0	82.2 ± 3.9	84.4 ± 4.2	83.0 ± 4.0
	Large	**0.051** ± 0.007	0.129 ± 0.020	0.146 ± 0.017	0.153 ± 0.018	0.134 ± 0.016	0.141 ± 0.017	117.2 ± 5.5	**22.8** ± 1.0	118.4 ± 5.5	116.7 ± 5.3	119.5 ± 5.8	117.8 ± 5.5
8	Small	**0.098** ± 0.009	0.261 ± 0.023	0.254 ± 0.024	0.262 ± 0.025	0.230 ± 0.022	0.245 ± 0.024	105.8 ± 4.9	**18.9** ± 0.9	106.8 ± 4.9	105.1 ± 4.7	108.1 ± 5.1	106.4 ± 4.8
	Medium	**0.108** ± 0.010	0.269 ± 0.025	0.266 ± 0.026	0.275 ± 0.027	0.248 ± 0.025	0.258 ± 0.026	157.4 ± 7.2	**24.5** ± 1.2	158.7 ± 7.2	156.4 ± 7.0	160.5 ± 7.5	158.1 ± 7.2
	Large	**0.118** ± 0.013	0.284 ± 0.031	0.292 ± 0.029	0.291 ± 0.031	0.275 ± 0.027	0.280 ± 0.028	215.3 ± 10.1	**32.1** ± 1.6	216.9 ± 10.1	214.2 ± 9.8	219.2 ± 10.5	216.5 ± 10.2

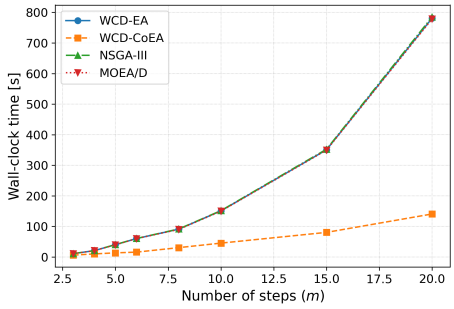


Fig. 7. Runtime scalability of all the tested methods versus the number of FlipIt game steps (m).

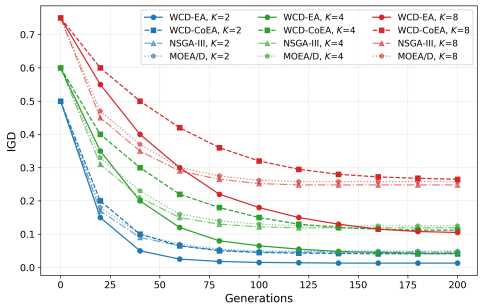


Fig. 8. Convergence analysis of WCD-EA, WCD-CoEA, and the set-based baseline variants of NSGA-III and MOEA/D for various number of payoff objectives (K) averaged over all game instances in the FlipIt security game.

point adaptations. To isolate the impact of the dominance relations, all baselines in this benchmark are equipped with the specialized mixture-merge crossover. The persistent performance gap between the WCD-EA and these alternative adaptations indicates that lacking the pessimistic worst-case formulation may discard critical structural information. Statistical comparisons between WCD-EA and all other tested methods using Welch's t-test revealed statistically significant differences in IGD scores ($p < 0.01$) across all benchmarks. This highlights the practical importance of the proposed worst-case dominance relation for effectively navigating the point-to-set mapping inherent to complex MO-SSGs.

The exact enumeration of Follower responses in WCD-EA allows more accurate estimation of the worst-case dominance relation. In contrast, WCD-CoEA relies on a low-fidelity surrogate objective that favors generalizable Follower responses but with the possibility of underestimating the Leader's worst-case outcome set. Our checks on the final populations reveal that the surrogate objective overestimates the Leader's performance by an average of 5% to 12% compared to the true worst-case outcome, confirming both the presence and magnitude of this approximation bias. This results in an IGD that is inferior to WCD-EA and closer to NSGA-III and MOEA/D. The approximate modeling of Follower decisions offers computational benefits, but at a noticeable reduction in accuracy.

2) Time Scalability: Figure 7 plots wall-clock time versus the number of game steps m. Because WCD-EA, NSGA-III, and MOEA/D require exact Follower response enumeration over full strategy space, their runtime curves exhibit a steeper slope. WCD-CoEA scales better, e.g., for $m = 20$ the WCD-EA runtime is about six times longer. For short horizons ($m \leq 6$) the time gap is small, but it expands steadily with m. For longer games ($m \geq 15$), WCD-CoEA provides a practical speed advantage, though without preserving an IGD similar to that achieved by WCD-EA (cf. Table II).

3) Convergence Analysis: Figure 8 illustrates the IGD trajectories over the course of evolution for $K \in \{2, 4, 8\}$. The WCD-EA quickly approaches low IGD values (e.g. ≈ 0.01 for $K = 2$). In contrast, WCD-CoEA exhibits slower initial convergence rate relative to all other algorithms, and only marginally surpasses the stagnating baseline MOEAs (NSGA-III and MOEA/D) towards the end of the evolutionary process. The standard MOEAs (NSGA-III and MOEA/D), despite a strong initial descent, prematurely stagnate at higher IGD values, reminiscent of their behavior in the ToW game.

As the number of objectives increases, the convergence rate naturally slows down, and the final IGD values remain higher, reflecting the inherent difficulty of covering high-dimensional objective spaces. Crucially, the performance gap between WCD-EA and WCD-CoEA becomes more pronounced as K increases. While the WCD-CoEA approximation is competitive in low-dimensional spaces, the separation widens significantly for $K = 8$. This suggests that the coevolutionary approximation of the Follower responses becomes less precise in higher dimensions compared to the exact enumeration used in WCD-EA, limiting the Leader's ability to fine-tune strategies against the worst-case scenarios.

Overall, the experiments illustrate a fundamental trade-off between fidelity of modeling the Follower and computational

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Transactions on Evolutionary Computation. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TEVC.2026.3749205

scalability. The coevolutionary mechanism establishes a clear trade-off: it sacrifices accuracy for substantial efficiency gains. WCD-CoEA yields a competitive wall-clock time, making it a preferable choice when the computational budget is constrained or many instances must be solved. In applications where the very best IGD is critical and runtime is less of a concern, WCD-EA would be a better option.

VIII. Conclusions

This work introduces a novel framework for *Multiobjective Stackelberg Security Games* (MO-SSGs), expanding classical single-objective models to settings where both Leader and Follower optimize multiple, potentially conflicting, objectives. We formalize key concepts such as the optimal set of Follower responses, the Leader's associated worst-case outcome set, worst-case dominance relations, and rationalizable Leader strategies, providing a theoretical foundation for multiobjective reasoning in Security Games. The proposed worst-case dominance relation induces a strict partial order over the Leader's strategy space, thereby yielding a natural ranking of strategies for evolutionary selection to operate on.

Given the structural incompatibility between worst-case dominance (under point-to-set mappings) and standard Pareto dominance (which is inherently point-to-point), two new algorithmic approaches specially crafted for MO-SSGs are put forward: the *Worst-Case Dominance Evolutionary Algorithm* (WCD-EA) based on exhaustive enumeration of optimal Follower responses, and the *Worst-Case Dominance Coevolutionary Algorithm* (WCD-CoEA) that jointly evolves Leader and Follower strategies. A central feature of these algorithms is a mixture–merge crossover operator that produces offspring while maintaining a compact representation, thereby facilitating the discovery of sparse mixed strategies. Experimental results on instructional Tug of War games (whose solutions can be analytically derived) and FlipIt games (inspired by real world cybersecurity interactions) demonstrate that the WCD-EA achieves the best solution quality, surpassing adaptations of existing MOEAs. In contrast, the WCD-CoEA offers substantial runtime savings, with competitive performance on some lower-dimensional problem instances. One limitation of the WCD-CoEA is that it employs a low-fidelity "surrogate" evaluation function that reduces computational overhead but at the cost of a potential bias: the Leader can underestimate the true severity of the Follower's worst-case response, leading to strategies that may be sub-optimal in a strict worst-case sense.

These current algorithmic trade-offs, alongside the practical requirements of real-world deployment, outline a clear trajectory for future research. Subsequent studies will focus on designing methods that better bridge the gap between the speed of coevolution and the rigorous accuracy of exhaustive enumeration. Transitioning this theoretical framework into operational security tools requires addressing the final *a posteriori* decision-making step. Just as in standard multiobjective optimization, a human decision-maker must ultimately select a single concrete strategy for execution from the returned set of rationalizable trade-offs. Developing dedicated post-hoc decision support techniques for this process is a natural next step to fully realize the practical potential of MO-SSGs.

References

[1] A. Sinha, T. H. Nguyen, D. Kar, M. Brown, M. Tambe, and A. X. Jiang, "From physical security to cybersecurity," *Journal of Cybersecurity*, vol. 1, no. 1, pp. 19–35, 2015.

[2] F. Fang, P. Stone, and M. Tambe, "When security games go green: Designing defender strategies to prevent poaching and illegal fishing." in *Proceedings of the 24th IJCAI Conference*, vol. 15, 2015, pp. 2589–2595.

[3] L.-A. Kirkland, A. de Waal, and J. P. de Villiers, "Evaluation of a pure-strategy Stackelberg game for wildlife security in a geospatial framework," in *Southern African Conference for Artificial Intelligence Research*. Springer, 2020, pp. 101–118.

[4] A. Sinha, F. Fang, B. An, C. Kiekintveld, and M. Tambe, "Stackelberg security games: looking beyond a decade of success," in *Proceedings of the 27th IJCAI Conference*, 2018, pp. 5494–5501.

[5] D. Kar, T. H. Nguyen, F. Fang, M. Brown, A. Sinha, M. Tambe, and A. X. Jiang, "Trends and applications in Stackelberg security games," *Handbook of dynamic game theory*, pp. 1–47, 2017.

[6] J. Pita, M. Jain, J. Marecki, F. Ordóñez, C. Portway, M. Tambe, C. Western, P. Paruchuri, and S. Kraus, "Deployed armor protection: the application of a game theoretic model for security at the los angeles international airport," in *Proceedings of the 7th AAMAS Conference*, 2008, pp. 125–132.

[7] J. Pita, M. Tambe, C. Kiekintveld, S. Cullen, and E. Steigerwald, "Guards—innovative application of game theory for national airport security," in *Proceedings of the 20th IJCAI Conference*, 2011, p. 2710.

[8] E. Shieh, B. An, R. Yang, M. Tambe, C. Baldwin, J. DiRenzo, B. Maule, and G. Meyer, "Protect: A deployed game theoretic system to protect the ports of the united states," in *Proceedings of the 11th AAMAS Conference*, 2012, pp. 13–20.

[9] J. Karwowski, J. Mańdziuk, A. Żychowski, F. Grajek, and B. An, "A memetic approach for sequential security games on a plane with moving targets," in *Proceedings of the 33rd AAAI Conference*, vol. 33, 2019, pp. 970–977.

[10] A. Żychowski, J. Mańdziuk, E. Bondi, A. Venugopal, M. Tambe, and B. Ravindran, "Evolutionary approach to Security Games with signaling," in *Proceedings of the 31st IJCAI Conference*, 2022, pp. 620–627.

[11] F. Fang, T. H. Nguyen, R. Pickles, W. Y. Lam, G. R. Clements, B. An, A. Singh, M. Tambe, and A. Lemieux, "Deploying paws: Field optimization of the protection assistant for wildlife security," in *Proceedings of the 28th Innovative Applications of Artificial Intelligence Conference*, 2016, pp. 3966–3973.

[12] M. Brown, B. An, C. Kiekintveld, F. Ordóñez, and M. Tambe, "An extended study on multi-objective security games," *Proceedings of the 13th AAMAS Conference*, vol. 28, no. 1, pp. 31–71, 2014.

[13] H. Qian, Y.-P. Wu, R.-J. Qin, X. An, Y. Chen, and A. Zhou, "Provable space discretization based evolutionary search for scalable multi-objective security games," *Swarm and Evolutionary Computation*, vol. 92, p. 101770, 2025.

[14] A. Żychowski, A. Gupta, J. Mańdziuk, and Y. S. Ong, "Addressing expensive multi-objective games with postponed preference articulation via memetic co-evolution," *Knowledge-Based Systems*, vol. 154, pp. 17–31, 2018.

[15] M. Harel, E. Eisenstadt-Matalon, and A. Moshaiov, "Solving zero-sum multi-objective games with a-priori secondary criteria," *Journal of Multi-Criteria Decision Analysis*, vol. 30, no. 1-2, pp. 3–23, 2023.

[16] J.-A. Mejía-De-Dios, A. Rodríguez-Molina, and E. Mezura-Montes, "Multiobjective bilevel optimization: A survey of the state-of-the-art," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 53, no. 9, pp. 5478–5490, 2023.

[17] H. Jiang, K. Chou, Y. Tian, X. Zhang, and Y. Jin, "Efficient surrogate modeling method for evolutionary algorithm to solve bilevel optimization problems," *IEEE Transactions on Cybernetics*, vol. 54, no. 7, pp. 4335–4347, 2023.

[18] G. Leitmann, "On generalized Stackelberg strategies," *Journal of Optimization Theory and Applications*, vol. 26, no. 4, pp. 637–643, 1978.

[19] B. Stengel and S. Zamir, "Leadership with commitment to mixed strategies," *CDAM Research Report LSE-CDAM-2004-01*, 2004.

[20] P. Paruchuri, J. P. Pearce, J. Marecki, M. Tambe, F. Ordonez, and S. Kraus, "Playing games for security: an efficient exact algorithm for solving bayesian Stackelberg games," in *Proceedings of the 7th AAMAS Conference*, 2008, pp. 895–902.

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Transactions on Evolutionary Computation. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TEVC.2026.3749205

[21] B. Bošanský and J. Čermak, "Sequence-form algorithm for computing Stackelberg equilibria in extensive-form games," in *Proceedings of the 29th AAAI Conference*, 2015, pp. 805–811.

[22] J. Čermák, B. Bošanský, K. Durkota, V. Lisý, and C. Kiekintveld, "Using correlated strategies for computing Stackelberg equilibria in extensive-form games," in *Proceedings of the 30th AAAI Conference*, 2016, pp. 439–445.

[23] J. Karwowski and J. Mańdziuk, "Stackelberg Equilibrium Approximation in General-Sum Extensive-Form Games with Double-Oracle Sampling Method," in *Proceedings of the 18th AAMAS Conference*, 2019, pp. 2045–2047.

[24] ——, "Double-oracle sampling method for Stackelberg Equilibrium approximation in general-sum extensive-form games," in *Proceedings of the 34th AAAI Conference*, vol. 34, 2020, pp. 2054–2061.

[25] M. Świechowski, K. Godlewski, B. Sawicki, and J. Mańdziuk, "Monte carlo tree search: A review of recent modifications and applications," *Artificial Intelligence Review*, vol. 56, no. 3, pp. 2497–2562, 2023.

[26] A. Żychowski and J. Mańdziuk, "A Generic Metaheuristic Approach to Sequential Security Games," in *Proceedings of the 19th AAMAS Conference*, 2020, p. 2089–2091.

[27] A. Żychowski and J. Mańdziuk, "Evolution of Strategies in Sequential Security Games," in *Proceedings of the 20th AAMAS Conference*, 2021, pp. 1434–1442.

[28] ——, "Coevolution of players strategies in security games," *Journal of Computational Science*, vol. 68, p. 101980, 2023.

[29] A. Żychowski, A. Gupta, Y. S. Ong, and J. Mańdziuk, "Augmented decision spaces for Stackelberg security games: Sparse evolution begets scalability," in *Proceedings of the GECCO Conference*, 2025, pp. 854–862.

[30] ——, "Scalability via sparsity in Stackelberg security games: An augmented decision space approach," *ACM Transactions on Evolutionary Learning*, 2026.

[31] E. Bondi, H. Oh, H. Xu, F. Fang, B. Dilkina, and M. Tambe, "To signal or not to signal: Exploiting uncertain real-time information in signaling games for security and sustainability," in *Proceedings of the 34th AAAI Conference*, 2020, pp. 1369–1377.

[32] T. H. Nguyen, R. Yang, A. Azaria, S. Kraus, and M. Tambe, "Analyzing the effectiveness of adversary modeling in security games," in *Proceedings of the 27th AAAI Conference*, 2013, pp. 718–724.

[33] R. Yang, C. Kiekintveld, F. Ordóñez, M. Tambe, and R. John, "Improving resource allocation strategies against human adversaries in security games: An extended study," *Artificial Intelligence*, vol. 195, no. 195, pp. 440–469, 2013.

[34] J. Pita, M. Jain, F. Ordóñez, M. Tambe, S. Kraus, R. Magori-Cohen, and M. Tambe, "Effective solutions for real-world Stackelberg games: When agents must deal with human uncertainties," *Security and Game Theory*, pp. 193–212, 2011.

[35] A. Żychowski and J. Mańdziuk, "Learning attacker's bounded rationality model in security games," in *Proceedings of the 28th ICONIP Conference*, vol. CCIS 1516, 2021, pp. 530–539.

[36] J. Karwowski, J. Mańdziuk, and A. Żychowski, "Sequential Stackelberg games with bounded rationality," *Applied Soft Computing*, vol. 132, p. 109846, 2023.

[37] M. Jain, E. Kardes, C. Kiekintveld, F. Ordónez, and M. Tambe, "Security games with arbitrary schedules: A branch and price approach," in *Proceedings of the 21st AAAI Conference*, vol. 24, no. 1, 2010, pp. 792–797.

[38] F. Ordónez, M. Tambe, J. F. Jara, M. Jain, C. Kiekintveld, and J. Tsai, "Deployed security games for patrol planning," in *Handbook of Operations Research for Homeland Security*. Springer, 2012, pp. 45–72.

[39] Y. Wang, Z. R. Shi, L. Yu, Y. Wu, R. Singh, L. Joppa, and F. Fang, "Deep reinforcement learning for green security games with real-time information," in *Proceedings of the 33rd AAAI Conference*, vol. 33, no. 01, 2019, pp. 1401–1408.

[40] K. Durkota, V. Lisỳ, B. Bosanskỳ, and C. Kiekintveld, "Optimal network security hardening using attack graph games." in *Proceedings of the 24th IJCAI Conference*, 2015, pp. 526–532.

[41] M. Jain, J. Tsai, J. Pita, C. Kiekintveld, S. Rathi, M. Tambe, and F. Ordónez, "Software assistants for randomized patrol planning for the lax airport police and the federal air marshal service," *Interfaces*, vol. 40, no. 4, pp. 267–290, 2010.

[42] E. Zitzler, L. Thiele, and J. Bader, "On set-based multiobjective optimization," *IEEE Transactions on Evolutionary Computation*, vol. 14, no. 1, pp. 58–79, 2009.

[43] C. A. C. Coello, G. B. Lamont, and D. A. V. Veldhuizen, *Evolutionary algorithms for solving multi-objective problems*. Springer, 2007.

[44] A. Sinha, P. Malo, and K. Deb, "A review on bilevel optimization: From classical to evolutionary approaches and applications," *IEEE Transactions on Evolutionary Computation*, vol. 22, no. 2, pp. 276–295, 2017.

[45] G. Eichfelder, "Multiobjective bilevel optimization," *Mathematical Programming*, vol. 123, no. 2, pp. 419–449, 2010.

[46] E. Zitzler, L. Thiele, M. Laumanns, C. M. Fonseca, and V. G. Da Fonseca, "Performance assessment of multiobjective optimizers: An analysis and review," *IEEE Transactions on evolutionary computation*, vol. 7, no. 2, pp. 117–132, 2003.

[47] D. Blackwell, "An analog of the minimax theorem for vector payoffs," *Pacific Journal of Mathematics*, vol. 6, pp. 1–8, 1956.

[48] B. Contini, "A decision model under uncertainty with multiple objectives," *Theory of Games: Techniques and Applications. New York*, vol. 9, 1966.

[49] D. Lozovanu, D. Solomon, and A. Zelikovsky, "Multiobjective games and determining pareto-nash equilibria," *Buletinul Academiei de Ştiinţe a Republicii Moldova. Matematica*, vol. 49, no. 3, pp. 115–122, 2005.

[50] M. S. Monfared, S. E. Monabbati, and A. R. Kafshgar, "Pareto-optimal equilibrium points in non-cooperative multi-objective optimization problems," *Expert Systems with Applications*, vol. 178, p. 114995, 2021.

[51] R. Rădulescu, P. Mannion, D. M. Roijers, and A. Nowé, "Multi-objective multi-agent decision making: a utility-based analysis and survey," *Autonomous Agents and Multi-Agent Systems*, vol. 34, no. 1, p. 10, 2020.

[52] E. Eisenstadt, A. Moshaiov, G. Avigad, and J. Branke, "Rationalizable strategies in multi-objective games under undecided objective preferences," in *Technical Report*. School of Mechanical Engineering, Tel-Aviv University, 2016.

[53] E. Eisenstadt, A. Moshaiov, and G. Avigad, "Co-evolution of strategies for multi-objective games under postponed objective preferences," in *2015 IEEE Conference on Computational Intelligence and Games (CIG)*. IEEE, 2015, pp. 461–468.

[54] E. Eisenstadt-Matalon and A. Moshaiov, "Decision analysis of rationalizable strategies in non-zero-sum multi-payoff games," *Intelligent Decision Technologies*, vol. 16, no. 3, pp. 487–504, 2022.

[55] A. Suresh, J. Kongmanee, K. Deb, and V. N. Boddeti, "Multi-objective coevolution and decision-making for cooperative and competitive environments," in *2021 IEEE Congress on Evolutionary Computation (CEC)*. IEEE, 2021, pp. 1601–1608.

[56] M. Breton, A. Alj, and A. Haurie, "Sequential Stackelberg equilibria in two-person games," *Journal of Optimization Theory and Applications*, vol. 59, no. 1, p. 71–97, 1988.

[57] A. Gupta and Y. S. Ong, "An evolutionary algorithm with adaptive scalarization for multiobjective bilevel programs," in *2015 IEEE Congress on Evolutionary Computation (CEC)*. IEEE, May 2015, pp. 1636–1642.

[58] K. Deb and M. Ehrgott, "On generalized dominance structures for multiobjective optimization," *Mathematical and Computational Applications*, vol. 28, no. 5, p. 100, 2023.

[59] V. Conitzer and T. Sandholm, "Computing the optimal strategy to commit to," in *Proceedings of the 7th ACM Conference on Electronic Commerce*, 2006, pp. 82–90.

[60] V. Gabrel, C. Murat, and A. Thiele, "Recent advances in robust optimization: An overview," *European Journal of Operational Research*, vol. 235, no. 3, pp. 471–483, 2014.

[61] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: NSGA-II," *IEEE Transactions on Evolutionary Computation*, vol. 6, no. 2, pp. 182–197, 2002.

[62] K. Deb and H. Jain, "An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, part i: solving problems with box constraints," *IEEE Transactions on Evolutionary Computation*, vol. 18, no. 4, pp. 577–601, 2013.

[63] Q. Zhang and H. Li, "MOEA/D: A multiobjective evolutionary algorithm based on decomposition," *IEEE Transactions on Evolutionary Computation*, vol. 11, no. 6, pp. 712–731, 2007.

[64] J. Blank and K. Deb, "Pymoo: Multi-objective optimization in Python," *IEEE Access*, vol. 8, pp. 89 497–89 509, 2020.

[65] M. Van Dijk, A. Juels, A. Oprea, and R. L. Rivest, "Flipit: The game of "stealthy takeover"," *Journal of Cryptology*, vol. 26, no. 4, pp. 655–713, 2013.

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/